PICA: Physics-Guided Implicit Neural Fields with Boundary Conditioning for Intracranial Aneurysm Hemodynamics

Mengfan Hu[1], Siyu Li[1], and Zeng Zhang[1]

School of Biomedical Engineering and State Key Laboratory of Advanced Medical Materials and Devices, ShanghaiTech University, Shanghai 201210, China
zhangzeng@shanghaitech.edu.cn

Abstract. Patient-specific hemodynamic assessment is critical for intracranial aneurysm rupture risk stratification, yet standard computational fluid dynamics (CFD) pipelines are too time-consuming for routine clinical use. Existing learning-based surrogates enable fast inference but often produce over-smoothed velocity fields, obscuring near-wall gradients and fine-scale vortical structures that are important for clinically relevant biomarkers. We propose PICA, a boundary-conditioned implicit neural field for continuous 3D hemodynamic prediction directly from voxelized aneurysm geometries. PICA encodes signed distance field (SDF)-augmented inputs with a multi-scale backbone and conditions the representation on patient-specific inlet parameters via feature-wise modulation, enabling robust adaptation across flow regimes. An implicit coordinate-based decoder queries the learned representation at arbitrary 3D locations to generate dense velocity fields without meshing or geometric reconstruction. To improve physical fidelity, we incorporate physics-guided regularization that enforces incompressibility and steady-state Navier–Stokes consistency, together with a progressive two-stage training curriculum to enhance cross-anatomy generalization. On the Aneumo benchmark, PICA achieves sub-second inference and reduces Mean Normalised Absolute Error by 25.4% compared to DeepONet-SwinT, while better recovering complex intra-saccular vortices and near-wall flow patterns.

Keywords: Intracranial Aneurysm · Hemodynamics · Implicit Neural Representations · Physics-Guided Deep Learning

1 Introduction

Intracranial aneurysms (IAs) affect 3 to 5% of the adult population and are a leading cause of aneurysmal subarachnoid hemorrhage, carrying high early mortality and severe neurological sequelae in survivors [4, 19]. Key hemodynamic parameters, such as wall shear stress (WSS), oscillatory shear index (OSI), and intra-saccular flow complexity, directly drive endothelial mechanobiology and progressive sac growth [6]. This interplay between biomechanical forces and wall degeneration makes patient-specific hemodynamic evaluation highly discriminative for rupture risk stratification [18, 19, 21]. Although computational fluid dynamics (CFD) is the gold standard for quantifying these forces, its conventional pipeline (vascular segmentation, mesh generation, and iterative Navier-Stokes solving) demands expert intervention and hours of computation per case [1, 8]. Consequently, CFD remains impractical for emergency triage, population screening, and intraoperative planning.

To address these computational bottlenecks, deep learning surrogates have emerged to approximate 3D blood flow fields at a fraction of the CFD cost [7,16]. However, voxel-based CNNs are limited by fixed spatial resolutions and poor near-wall flow fidelity. Operator-learning methods like Deep Operator Networks (DeepONets) [3, 11] and Fourier Neural Operators (FNOs) [5, 9] offer mesh-free inference but suffer from systematic over-smoothing. This artifact suppresses crucial vortical structures and near-wall low-WSS gradients, which are precisely the features correlated with pathological wall degradation and imminent rupture [19]. Implicit neural representations (INRs) [12, 13] excel in coordinate-based continuous signal encoding yet remain largely unexplored for aneurysm flow modeling. The recent Aneumo benchmark [8] highlights these inherent limitations in existing DeepONet and Swin Transformer baselines, underscoring the urgent need for a framework that ensures physical consistency, boundary condition coupling, and high-fidelity local flow resolution.

To bridge this gap, we propose **PICA** (**P**hysics-guided **I**mplicit **C**onditional **A**neurysm hemodynamics network). PICA recasts 3D hemodynamic prediction as learning a boundary condition-guided implicit neural field over an SDF-enriched voxel representation [14]. A multi-scale 3D U-Net encoder [2] extracts hierarchical geometric features, while Feature-wise Linear Modulation (FiLM) [15] explicitly conditions these features on patient-specific boundary parameters (inlet centroid position, surface normal direction, and volumetric flow rate Q). A coordinate-aware implicit decoder with Fourier positional encoding [20] then generates continuous velocity predictions at arbitrary 3D coordinates. To guarantee physical fidelity without requiring extra labeled data, dual physics-guided regularization terms enforce divergence-free and momentum-conserving outputs [17].

In summary, our main contributions are:

- We introduce PICA, an INR-based surrogate integrating SDF-enriched geometry, hierarchical FiLM boundary conditioning, and physics-guided regularization for mesh-free 3D aneurysm hemodynamic prediction.
- By enforcing divergence-free and vorticity transport constraints, PICA substantially improves the near-wall velocity field and recovers fine-grained vortical structures suppressed by existing operator-learning surrogates.
- PICA reduces MNAE by 25.4% over DeepONet-SwinT while maintaining sub-second inference, establishing state-of-the-art performance on the Aneumo benchmark.

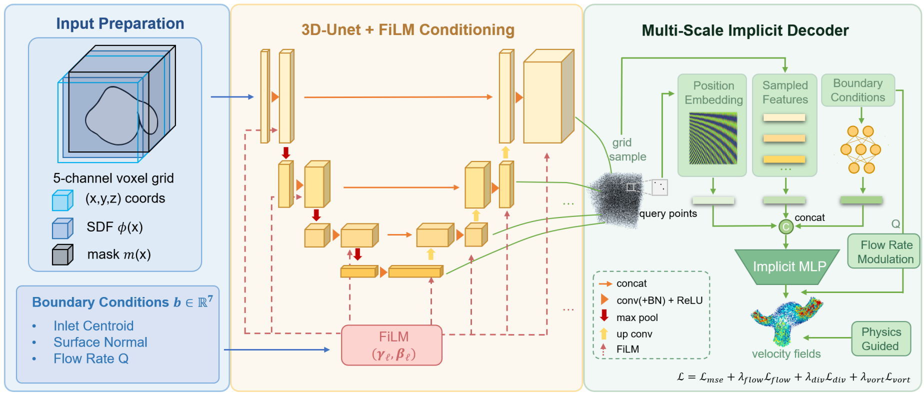

Fig. 1. Overview of the proposed PICA framework. **Left:** The input preparation module ingests a 5-channel voxel grid (3D spatial coordinates, SDF $\phi(\mathbf{x})$, and binary mask $m(\mathbf{x})$) alongside patient-specific boundary conditions $\mathbf{b} \in \mathbb{R}^7$. **Middle:** A 3D U-Net extracts multi-scale geometric features, while FiLM layers dynamically inject boundary parameters (inlet centroid, surface normal, and flow rate Q) across the feature hierarchy. **Right:** The Multi-Scale Implicit Decoder queries features at arbitrary 3D coordinates via trilinear interpolation, concatenates Fourier positional embeddings and boundary conditions, and predicts continuous 3D velocity fields supervised by a composite loss ($\mathcal{L}_{\text{mse}}$, $\mathcal{L}_{\text{div}}$, $\mathcal{L}_{\text{vort}}$) enforcing physical consistency.

2 Method

PICA maps a five-channel voxelized geometric representation $\mathcal{V} \in \mathbb{R}^{5\times128^3}$ and patient-specific inlet boundary conditions $\mathbf{b} \in \mathbb{R}^7$ (centroid, normal direction, and flow rate Q) to a continuous 3D velocity field $\hat{\mathbf{u}}(\mathbf{x}) \in \mathbb{R}^3$ at arbitrary query coordinates $\mathbf{x}$. As illustrated in Fig. 1, the framework comprises an SDF-augmented encoder with FiLM conditioning, a multi-scale implicit decoder, and physics-guided regularization.

2.1 Geometric Encoding and FiLM Conditioning

Input Representation. While the binary lumen mask $m(\mathbf{x})$ delineates the fluid domain, the additional explicit spatial coordinates (x, y, z) and SDF $\phi(\mathbf{x})$ provide absolute positional context and continuous wall-proximity awareness. This enriched representation is critical for resolving steep near-wall velocity gradients that a mask-only input cannot capture, as confirmed by our ablation study (Table 2).

Boundary Vector Construction. For each case, the boundary vector is defined as $\mathbf{b} = [c_x, c_y, c_z, n_x, n_y, n_z, Q] \in \mathbb{R}^7$, where $\mathbf{c}$ is the inlet centroid, $\mathbf{n}$ is the unit inlet normal, and Q is the prescribed volumetric flow rate. The centroid and normal are computed automatically from the inlet boundary surface during preprocessing, without using the target CFD velocity or pressure fields; Q is taken from the simulation boundary-condition metadata and normalized using training-set statistics.

Boundary Conditioning via FiLM. $\mathcal{V}$ is processed by a symmetric 3D U-Net backbone. To adapt geometric features to patient-specific hemodynamics, boundary conditions $\mathbf{b}$ are injected across all encoder and decoder levels using Feature-wise Linear Modulation (FiLM) [15]. At each network stage ℓ, an MLP maps $\mathbf{b}$ into affine parameters $(\boldsymbol{\gamma}_\ell, \boldsymbol{\beta}_\ell)$, modulating the intermediate feature maps as $\widetilde{F}_\ell = \boldsymbol{\gamma}_\ell(\mathbf{b}) \odot F_\ell + \boldsymbol{\beta}_\ell(\mathbf{b})$. This mechanism effectively couples the static geometry with dynamic flow regimes throughout the feature hierarchy.

2.2 Coordinate-Aware Implicit Decoding

To predict continuous velocity fields without mesh generation, features are queried at arbitrary continuous coordinates $\mathbf{x}$. First, $\mathbf{x}$ is mapped to 63-dimensional high-frequency components via Fourier positional encoding [12] to capture fine-grained flow details. Concurrently, multi-resolution geometric features $\mathbf{f}(\mathbf{x})$ are sampled from four U-Net decoder stages via trilinear interpolation and fused using a learned weighted sum.

The implicit decoder, modeled as a 12-layer MLP, concatenates the Fourier-encoded coordinates, the aggregated geometric features $\mathbf{f}(\mathbf{x})$, and the re-encoded boundary conditions to predict base velocities $\hat{\mathbf{u}}_0 = [\hat{u}_0, \hat{v}_0, \hat{w}_0]^\top$. Finally, an auxiliary MLP branch predicts component-wise scaling factors $s_\eta(Q) \in \mathbb{R}^3_+$ derived from the inlet flow rate Q. The final physical velocity field is obtained via multiplicative modulation:

$$\hat{\mathbf{u}}(\mathbf{x}) = \big(s_\eta(Q) + \mathbf{1}\big) \odot \hat{\mathbf{u}}_0, \tag{1}$$

where $\mathbf{1} \in \mathbb{R}^3$ denotes the all-ones vector, whose inclusion as a learnable bias initializes the scale factors near unity to stabilize early training dynamics.

2.3 Physics-Guided Training Objective

PICA is supervised by a composite loss $\mathcal{L}$ evaluated over a set of query points $\mathcal{P}$, which combines data-driven fidelity with physics-guided regularization:

$$\mathcal{L} = \mathcal{L}_{\text{mse}} + \lambda_{\text{flow}}\mathcal{L}_{\text{flow}} + \lambda_{\text{div}}\mathcal{L}_{\text{div}} + \lambda_{\text{vort}}\mathcal{L}_{\text{vort}}. \tag{2}$$

To balance the dimensionally heterogeneous loss terms, we adopt an uncertainty-based adaptive weighting strategy. Each auxiliary objective is associated with a learnable log-variance parameter, allowing the model to automatically adjust the relative contribution of data fidelity, flow-rate consistency, incompressibility, and vorticity transport constraints during training. We initialize the auxiliary terms with scale factors $\lambda_{\text{flow}} = 10^{-6}$, $\lambda_{\text{div}} = 10^{-3}$, and $\lambda_{\text{vort}} = 10^{-4}$ for numerical stability, while the final effective weights are learned jointly with the network parameters.

Fidelity and Scaling. The primary data-driven objective is the mean squared error (MSE) between the predicted velocity $\hat{\mathbf{u}}$ and the CFD ground-truth $\mathbf{u}^*$:

$$\mathcal{L}_{\text{mse}} = \frac{1}{|\mathcal{P}|} \sum_{\mathbf{x} \in \mathcal{P}} \left\| \hat{\mathbf{u}}(\mathbf{x}) - \mathbf{u}^*(\mathbf{x}) \right\|^2. \tag{3}$$

Additionally, assuming velocity scales linearly with the inlet flow rate Q under steady viscous conditions, we apply a soft regularization prior that penalizes inconsistencies in flow-rate-normalized mean velocity magnitudes:

$$\mathcal{L}_{\text{flow}} = \left(\frac{1}{|\mathcal{P}|} \sum_{\mathbf{x} \in \mathcal{P}} \frac{\|\hat{\mathbf{u}}(\mathbf{x})\|}{Q} - \frac{1}{|\mathcal{P}|} \sum_{\mathbf{x} \in \mathcal{P}} \frac{\|\mathbf{u}^*(\mathbf{x})\|}{Q} \right)^2, \tag{4}$$

where the summations are taken over all query points $\mathbf{x} \in \mathcal{P}$ sampled within the lumen.

Incompressibility Loss. To enforce mass conservation for incompressible blood flow, we impose a soft divergence-free constraint. Spatial partial derivatives are computed analytically via automatic differentiation, penalizing the divergence residual:

$$\mathcal{L}_{\text{div}} = \frac{1}{|\mathcal{P}|} \sum_{\mathbf{x} \in \mathcal{P}} \left(\nabla \cdot \hat{\mathbf{u}} \right)^2. \tag{5}$$

Vorticity Transport Loss. Addressing the over-smoothing artifacts typical of deep learning surrogates requires explicit supervision of local flow rotational structures. Under steady-state assumptions, the vorticity $\boldsymbol{\omega} = \nabla \times \hat{\mathbf{u}}$ must satisfy the steady vorticity transport equation, derived by taking the curl of the Navier-Stokes momentum equation. We penalize this residual at each query point:

$$\mathcal{L}_{\text{vort}} = \frac{1}{|\mathcal{P}|} \sum_{\mathbf{x} \in \mathcal{P}} \left\| (\hat{\mathbf{u}} \cdot \nabla)\boldsymbol{\omega} - (\boldsymbol{\omega} \cdot \nabla)\hat{\mathbf{u}} - \nu \nabla^2 \boldsymbol{\omega} \right\|^2, \tag{6}$$

where ν denotes kinematic viscosity. This formulation explicitly models the material advection of vorticity against vortex stretching $(\boldsymbol{\omega} \cdot \nabla)\hat{\mathbf{u}}$ and viscous diffusion $\nu \nabla^2 \boldsymbol{\omega}$, effectively steering the network toward physically admissible, fine-grained solutions without requiring extra labeled data.

2.4 Two-Stage Progressive Training

To improve convergence and cross-anatomy generalization, we adopt a two-stage curriculum using identical training data for both stages.

Stage 1 (End-to-End Pretraining): The entire network is trained to establish stable, geometry-aware feature hierarchies that capture global vessel topology and local wall proximity.

Stage 2 (Decoder-Focused Fine-Tuning): The U-Net encoder weights are frozen, and only the FiLM generators and implicit decoder MLP are updated. Concentrating gradient updates on boundary-condition adaptation and local coordinate decoding prevents the degradation of generalized geometric features and mitigates overfitting.

3 Experiments and Results

3.1 Experimental Settings

Dataset. We evaluate PICA on the publicly available Aneumo benchmark [8], a large-scale repository of steady-state Navier-Stokes CFD simulations for patient-specific intracranial aneurysms. To ensure geometric independence and prevent shape leakage, the selection of the training and testing sets strictly follows the standard benchmark split established in the original paper. From the raw simulation data, we extract dense pointwise velocity ground-truth within the lumen, and construct a 5-channel voxel grid $\mathcal{V} \in \mathbb{R}^{5\times128^3}$ alongside a boundary vector $\mathbf{b} \in \mathbb{R}^7$ as our model inputs.

Implementation Details. Voxel grids are generated at an isotropic 128^3 resolution. During training, $M = 4,096$ query points are uniformly sampled within the lumen per forward pass. The PICA model ($\sim$5M parameters) is optimized using AdamW (learning rate 10^{-4}, weight decay 10^{-5}). Stage 1 pretrains all parameters end-to-end for 5000 epochs (batch size 4). Stage 2 freezes the encoder, fine-tuning only the FiLM layers and implicit decoder for an additional 1000 epochs. Models were trained on two NVIDIA RTX 6000 Ada GPUs via DistributedDataParallel.

Baselines and Metrics. We compare PICA against two representative neural operators from the Aneumo benchmark [8]: **DeepONet** [11], a branch-trunk architecture lacking explicit voxel geometry encoding; and **DeepONet-SwinT** [8], an enhanced variant using a Swin Transformer [10] branch to capture volumetric features. Both baselines were trained from scratch using the identical data split. We focus on these directly comparable benchmark baselines because they share the same data split, input-output setting, and evaluation protocol; broader comparisons with graph-based or continuous-field CFD surrogates are left for future work. We report the component-wise relative ℓ_2 errors (e_u, e_v, e_w), Mean Normalised Absolute Error (MNAE $= \frac{1}{N}\sum_{i=1}^{N} |y_i - \hat{y}_i|/(\max(\mathbf{y}) - \min(\mathbf{y}))$), and per-case inference time.

3.2 Experimental Results

Quantitative Performance. Table 1 presents the predictive performance on the test set. PICA achieves the lowest errors across all metrics. Compared to the strongest baseline, DeepONet-SwinT, PICA reduces the MNAE by 25.4% while maintaining comparable sub-second inference speeds. The significant improvement over plain DeepONet highlights that point-cloud or mesh queries alone are insufficient; integrating SDF-enriched voxel features with hierarchical FiLM conditioning is crucial for resolving complex near-wall gradients under limited training data.

Qualitative Analysis. Fig. 2 visualizes cross-sectional velocity magnitude contours alongside absolute error distributions. DeepONet broadly overestimates intra-saccular velocities due to its lack of spatial context. DeepONet-SwinT captures the general topology but distorts the central high-velocity jet. In contrast,

Table 1. Quantitative comparison on the Aneumo test set. e_u, e_v, e_w: component-wise relative ℓ_2 error (%). MNAE: Mean Normalised Absolute Error. Best results are highlighted in **bold**.

Method	e_u (↓)	e_v (↓)	e_w (↓)	MNAE (↓)	Time (s) (↓)
DeepONet [11]	0.5044	0.4930	0.4651	0.0635	**0.0020**
DeepONet-SwinT [8]	0.4480	0.4306	0.3990	0.0512	0.0100
PICA (Ours)	**0.3531**	**0.3770**	**0.3466**	**0.0382**	0.0083

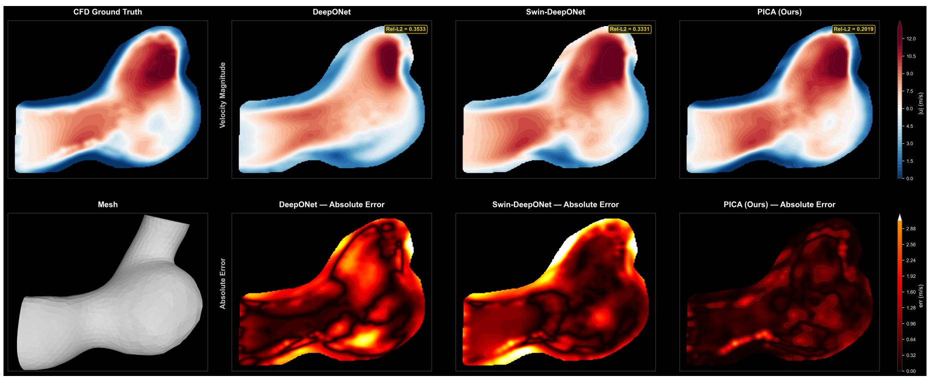

Fig. 2. Qualitative comparison of velocity magnitudes on a representative test geometry (Case 191). Top: Cross-sectional contour maps of CFD ground-truth and surrogate predictions, annotated with relative ℓ_2 errors. **Bottom:** 3D geometric mesh and corresponding absolute error maps. PICA significantly reduces local residuals, especially near the high-velocity jet and the vessel wall.

PICA recovers sharp boundary transitions and accurately delineates primary flow structures. The error maps (bottom row) confirm that PICA's residuals are sparse and low-magnitude, outperforming the baselines.

Further 3D flow evaluations are presented in Fig. 3. Streamline visualizations (Fig. 3a) reveal that both baselines suffer from severe over-smoothing, failing to reconstruct the complex swirling vortex within the aneurysm sac. PICA, guided by its vorticity transport loss, faithfully reproduces these fine-grained vortical structures, achieving an ℓ_2 error of 0.347. Furthermore, near-wall flow distributions (Fig. 3b) show that PICA accurately maps the high- and low-flow spatial regimes, effectively overcoming the systematic over-smoothing limitations of existing models.

3.3 Ablation Studies

Table 2 quantifies the contribution of PICA's core components using the test-set MNAE. Replacing the 5-channel input with a mask-only representation causes the largest performance drop (+52.1%), indicating that explicit spatial coordinates and SDF are the most critical cues for near-wall velocity resolution. Re-

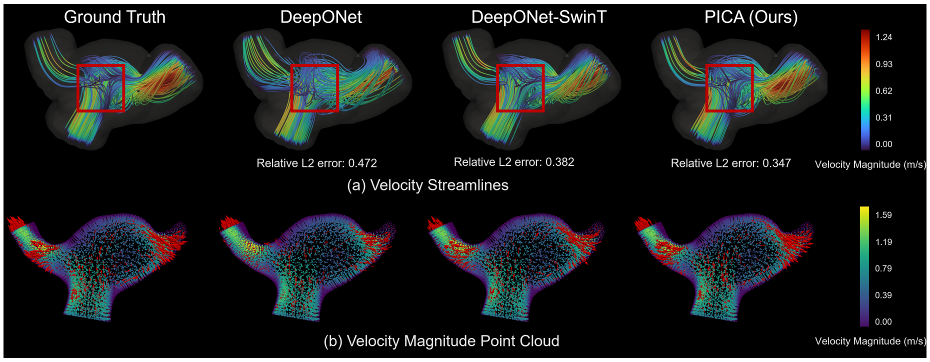

Fig. 3. 3D hemodynamic flow reconstruction. (a) Velocity streamlines highlighting complex intra-saccular recirculation (red boxes). PICA effectively mitigates the over-smoothing seen in baselines. **(b)** Cross-sectional velocity magnitude point clouds demonstrating PICA's superior fidelity in capturing near-wall gradients.

Table 2. Ablation study on the Aneumo test set. MNAE is averaged across all velocity components. Δ denotes the relative error increase compared to the full model ($\uparrow$ indicates a performance drop when the component is removed). Best result is in **bold**.

Configuration	**MNAE** ($\downarrow$)	Δ
w/o Geometric Encoding (Mask only)	0.0581	+52.1% ($\uparrow$)
w/o Hierarchical FiLM (Late fusion)	0.0493	+29.0% ($\uparrow$)
w/o Multi-scale Querying	0.0452	+18.3% ($\uparrow$)
w/o Physics Constraints ($\mathcal{L}_{\text{flow,div,vort}}$)	0.0438	+14.6% ($\uparrow$)
w/o Two-stage Training (Single-stage)	0.0415	+8.6% ($\uparrow$)
PICA (Full Model)	**0.0382**	–

ducing the 9-stage hierarchical FiLM to a single late-fusion concatenation raises errors by 29.0%, confirming the necessity of deep boundary-condition adaptation. Removing the multi-scale decoder feature sampling degrades accuracy by 18.3%, validating the importance of coarse-to-fine aggregation. Omitting the physics constraints ($\mathcal{L}_{\text{flow,div,vort}}$) impairs performance by 14.6%, specifically harming the prediction of high-velocity jets and vortical flows. Finally, reverting to single-stage training (without encoder freezing) increases errors by 8.6%, proving that our progressive curriculum actively mitigates overfitting and enhances geometric feature generalization without requiring additional data.

4 Discussion and Conclusion

In this work, we propose PICA, a physics-guided implicit neural framework that recasts 3D intracranial aneurysm hemodynamic prediction as a boundary-conditioned continuous field learning problem. By synergistically integrating SDF-enriched geometric encoding with hierarchical FiLM-based boundary conditioning, PICA circumvents the severe over-smoothing artifacts pervasive in existing operator-learning surrogates. Evaluated on the Aneumo benchmark, our method significantly outperforms state-of-the-art DeepONet baselines, recovering fine-grained near-wall gradients and complex vortical structures while maintaining sub-second inference speeds. Furthermore, our two-stage progressive training curriculum establishes highly robust geometric representations, ensuring strong cross-anatomy generalization under limited training data.

Despite these advances, PICA currently focuses on steady-state velocity predictions within synthetic geometries. This setting enables efficient reconstruction of patient-specific velocity fields and serves as a foundation for future extensions toward richer hemodynamic analyses. Temporal indices such as OSI and pressure-related analyses require time-resolved pulsatile modeling, and will therefore be explored in future extensions. Future work will further extend the implicit formulation to capture pulsatile flow dynamics and intravascular pressure variations, compare with broader surrogate families, and validate the framework on in-vivo, patient-specific clinical data.

Acknowledgments. This work is supported by ShanghaiTech AI4S Initiative SHTAI4S 202404.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Cebral, J.R., Mut, F., Weir, J., Putman, C.M.: Association of hemodynamic characteristics and cerebral aneurysm rupture. American Journal of Neuroradiology **32**(2), 264–270 (2011)
2. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: International conference on medical image computing and computer-assisted intervention. pp. 424–432. Springer (2016)
3. Cruz-González, O.L., Deplano, V., Ghattas, B.: Enhanced vascular flow simulations in aortic aneurysm via physics-informed neural networks and deep operator networks. Mechanics Research Communications **153**, 104642 (2026)
4. Etminan, N., Rinkel, G.J.: Unruptured intracranial aneurysms: development, rupture and preventive management. Nature Reviews Neurology **12**(12), 699–713 (2016)
5. Flouris, K., Halter, M., Lee, Y.Y., Castonguay, S., Jacobs, L., Dirix, P., Nestmann, J., Kozerke, S., Konukoglu, E.: Localized fno for spatiotemporal hemodynamic upsampling in aneurysm mri. In: International Workshop on Reconstruction and Imaging Motion Estimation. pp. 65–75. Springer (2025)
6. Kadem, M., Garber, L., Abdelkhalek, M., Al-Khazraji, B.K., Keshavarz-Motamed, Z.: Hemodynamic modeling, medical imaging, and machine learning and their applications to cardiovascular interventions. IEEE Reviews in Biomedical Engineering **16**, 403–423 (2022)

7. Kochkov, D., Smith, J.A., Alieva, A., Wang, Q., Brenner, M.P., Hoyer, S.: Machine learning–accelerated computational fluid dynamics. Proceedings of the National Academy of Sciences **118**(21), e2101784118 (2021)
8. Li, X., Zhou, Y., Xiao, F., Guo, X., Zhang, Y., Jiang, C., Ge, J., Wang, X., Wang, Q., Zhang, T., et al.: Aneumo: A large-scale comprehensive synthetic dataset of aneurysm hemodynamics. arXiv preprint arXiv:2501.09980 (2025)
9. Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., Anandkumar, A.: Fourier neural operator for parametric partial differential equations. arXiv preprint arXiv:2010.08895 (2020)
10. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
11. Lu, L., Jin, P., Pang, G., Zhang, Z., Karniadakis, G.E.: Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. Nature machine intelligence **3**(3), 218–229 (2021)
12. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM **65**(1), 99–106 (2021)
13. Molaei, A., Aminimehr, A., Tavakoli, A., Kazerouni, A., Azad, B., Azad, R., Merhof, D.: Implicit neural representation in medical imaging: A comparative survey. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2381–2391 (2023)
14. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 165–174 (2019)
15. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
16. Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., Battaglia, P.: Learning mesh-based simulation with graph networks. In: International conference on learning representations (2020)
17. Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational physics **378**, 686–707 (2019)
18. Sforza, D.M., Putman, C.M., Cebral, J.R.: Hemodynamics of cerebral aneurysms. Annual review of fluid mechanics **41**(1), 91–107 (2009)
19. Soldozy, S., Norat, P., Elsarrag, M., Chatrath, A., Costello, J.S., Sokolowski, J.D., Tvrdik, P., Kalani, M.Y.S., Park, M.S.: The biophysical role of hemodynamics in the pathogenesis of cerebral aneurysm formation and rupture. Neurosurgical focus **47**(1), E11 (2019)
20. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. Advances in neural information processing systems **33**, 7537–7547 (2020)
21. Turjman, A.S., Turjman, F., Edelman, E.R.: Role of fluid dynamics and inflammation in intracranial aneurysm formation. Circulation **129**(3), 373–382 (2014)